

Applying causality principles to the axiomatization of probabilistic cellular automata

Pablo Arrighi^{1,2}, Renan Fargetton², Vincent Nesme³, and Eric Thierry¹

¹ École normale supérieure de Lyon, LIP, 46 allée d’Italie, 69008 Lyon, France
 {parrighi,ethierry}@ens-lyon.fr

² Université de Grenoble, LIG, 220 rue de la chimie, 38400 SMH, France
 renan.fargetton@imag.fr

³ Universität Potsdam, Karl-Liebknecht-Str. 24/25, 14476 Potsdam, Germany
 nesme@qipc.org

Abstract. Cellular automata (CA) consist of an array of identical cells, each of which may take one of a finite number of possible states. The entire array evolves in discrete time steps by iterating a global evolution G . Further, this global evolution G is required to be shift-invariant (it acts the same everywhere) and causal (information cannot be transmitted faster than some fixed number of cells per time step). At least in the classical [13], reversible [17] and quantum cases [1], these two top-down axiomatic conditions are sufficient to entail more bottom-up, operational descriptions of G . We investigate whether the same is true in the probabilistic case.

Keywords: Characterization, noise, Markov process, stochastic Einstein locality, screening-off, common cause principle, non-signalling, Multi-party non-local box.

1 Introduction

Due to its built-in symmetries, CA constitute a clearly physics-like model of computation [18]. They model spatially distributed computation in space as we know it [25, 19], and therefore they constitute a framework for studying and proving properties about such systems – by far the most established. Conceived amongst theoretical physicists such as Ulam and Von Neumann [29], CA were soon considered as a possible model for particle physics, fluids, and the related differential equations. There are numerous results on this approach, often under the name of Lattice-gas cellular automata [33, 32, 9, 24]. More generally, CA are one of the main theoretical tools of ‘complex sciences’, where one studies the emergence of complex, global behaviours as arising from simple local interactions. There, CA have proved useful for modelling an incredible variety of things ranging traffic jams [20] to demographics and regional development or consumption [6, 31].

Each of this variety of contexts brings its own set of reasons to study probabilistic extensions of CA. This trend of work has already started, and generated

fascinating questions. For example, are there CA which can defend themselves against an everywhere present noise? Or is it the case that any initial configurations will ultimately be entirely erased, i.e. that such CA have only one limit distribution? In [26, 27] it was shown, via an extremely convoluted argument which many would like to simplify [23, 7, 12], that there exists CA which are resistant to this noise model. However, perhaps it would help to address these issues to have a robust definition of Probabilistic CA (PCA). All of these papers

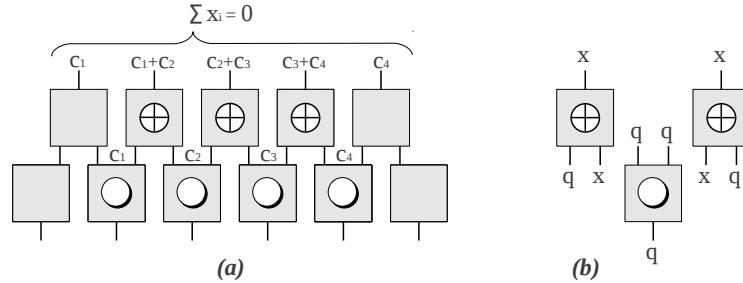


Fig. 1. (a) *Parity*. Time flows upwards. Wires are bits. Boxes are (randomized) gates, i.e. stochastic matrices. Boxes marked with a circle are coin tosses. Boxes marked \oplus perform addition modulo two. Input cells $0 \dots n$ are ignored. Output random variables $(X_i)_{i=1 \dots n}$ have the property that for strict subset I of $1 \dots n$, X_I is uniformly distributed, yet their global parity is always 0. **(b)** As such *Parity* is a localized stochastic map, but it can be modified into a translation-invariant stochastic map over entire configurations of the form $\dots qq01001qq\dots$, by ensuring that for any $x \in \{q, 0, 1\}$, $qx \mapsto x$, $xq \mapsto x$, $qq \mapsto q$, and hence making the coin tosses conditional upon their input cells not being q . *Parity* does not pertain to the class of Standard-PCA, but it ought to be considered a valid PCA.

begin with a definition of PCA, but these definitions are all variations of the same concept. Namely, these are stochastic maps having the property that they decompose in two phases: first, a classical CA is applied, and second, a model of noise is applied (a stochastic matrix is applied homogeneously on each individual cell). We refer to this class of stochastic maps as Standard-PCA and now explain its several drawbacks.

First, this class is incomplete. There are many stochastic maps which ought to be called PCA but are not Standard-PCA, not even with increased cell or neighbourhood sizes. *Parity* is a concrete example of this, see Figure 1. Indeed, *Parity* is translation-invariant, implementable by local means, but it may generate statistical correlations between any two adjacent regions of cells, which is never the case of Standard-PCA. Of course *Parity* can be implemented as the square (i.e. two time-steps) of a Standard-PCA, which points out to a second, even worse problem with this class: the composition of two Standard-PCA is not necessarily a Standard-PCA. This is counter-intuitive for a notion of PCA. Thirdly, the intuition behind this class is simple but ad-hoc; i.e. not founded

upon some meaningful high-level principles.

This paper aims to define PCA in a more robust manner, which is a long-standing problem. Criteria for a good axiomatic definition include: being composable, and being based on high-level principles, while entailing an operational description (i.e. implementation by local mechanisms).

Why not start right-away from the latter, operational description? If not entailed by well-agreed principles, the operational description may be incomplete and ad-hoc, as with Standard-PCA. Moreover, axiomatic definitions tend to have a practical interest as simple characterizations of the operational descriptions. In fact a great deal of foundational mathematical results work that way. When an axiomatic definition is given an operational description, we speak of a ‘structure theorem’ or a ‘representation theorem’ (e.g. spectral decomposition of a unitary operator as $\sum e^{i\theta_i} |\psi_i\rangle\langle\psi_i|$). When an operational description is given an axiomatic definition, we talk of a ‘characterization’ (e.g. Hedlund’s characterization of a CA as the set of shift-invariant continuous functions in the Cantor topology).

Why not generalize Hedlund’s characterization of CA to PCA? We will begin by investigating precisely that route and show that continuity arguments fail to characterize PCA (Section 3), as they do not forbid spontaneous generation of statistical correlation in arbitrarily distant places (Section 4). Counter-examples are necessarily non-signalling, non-local stochastic maps. That such objects can exist is now a well-known fact in quantum theory, where non-signalling non-locality arises from entanglement in Bell’s test experiment [5]. Recently their study has been abstracted away from quantum theory, for instance via the rather topical *NLBox* [4, 21, 3].

This points out the weakness of the non-signalling condition in the probabilistic/stochastic setting, which is a well-known issue in foundations of physics. In this context, more robust causality principles have been considered. In fact, Bell’s test experiments are motivated by a ‘local causality’, of which there exist several variants, all of them stemming from Reichenbach’s ‘principle of common cause’. Now, since PCA are nothing but simplifications of those spaces used in physics, this principle of common cause, if it is at all robust, should therefore lead to the axiomatization of PCA. We investigate this intriguing question (Section 5) and answer it in an unexpected fashion (Section 6). First we formalize the problem.

2 Problem statement

Definition 1 (Configurations) *A finite (unbounded) configuration c over a finite alphabet Σ is a function $c : \mathbb{Z} \rightarrow \Sigma$, with $i \mapsto c(i) = c_i$, such that there exists a (possibly empty) finite set I satisfying $i \notin I \Rightarrow c_i = q$, where q is a distinguished quiescent state of Σ . We denote \mathcal{C} , this set of finite configurations; it is countable.*

Countability is not crucial to the arguments developed in this paper, but it certainly makes the following definitions much easier:

Definition 2 (Random variables, states) Random variables denoted by X range over \mathcal{C} , i.e. entire configurations. Random variables denoted by X_i range over Σ , i.e. cells. Random variables denoted by X_I range over $I \rightarrow \Sigma$, i.e. sets of cells. Random variables corresponding to time t , are denoted X^t , X_I^t . The state of the random variable Y , denoted ρ_Y , is the functions from $\text{range}(Y)$ to $[0, 1]$ such that $\rho_Y(y) = \Pr(Y = y)$, with $\Pr(Y = y)$ the probability mass function of Y at y , i.e. they denote the law of distribution. For convenience in the particular case when Y is fully determined on y , i.e. such that $\rho_Y(x) = \delta_{xy}$, the state ρ_Y will be written \tilde{y} . Moreover for convenience $\rho_{X_I^t}$ will be denoted ρ_I^t , and referred to as the state of cells I at time t . Whenever I is finite, we can make ρ_I^t explicit as the vector $(\rho_I^t(w))_{w \in \Sigma^I}$ with the $|\Sigma^I|$ entries listed in the lexicographic order.

Definition 3 (Stochastic maps) Consider X and X' two random variables over the same range, and their corresponding states. We define the state $p\rho_X + (1-p)\rho_{X'}$ so that $(p\rho_X + (1-p)\rho_{X'})(y) = p\rho_X(y) + (1-p)\rho_{X'}(y)$. Consider S a function from states ($\text{range}(X) \rightarrow [0, 1]$) to states ($\text{range}(Y) \rightarrow [0, 1]$). Then S is a stochastic map over the range of X if and only if it is linear, i.e. $S(p\rho_X + (1-p)\rho_{X'}) = pS\rho_X + (1-p)S\rho_{X'}$. Whenever $\text{range}(X)$ and $\text{range}(Y)$ are finite, we can make S explicit as the stochastic matrix $S = (S\tilde{x}(y))_{y \in \text{range}(Y), x \in \text{range}(X)}$. Notice that each column is the law of distribution $S\tilde{x}$ for some state \tilde{x} , hence its entries are in $[0, 1]$ and sum to one (left stochasticity). We assume that the random variables over configurations $(X^t)_{t \in \mathbb{N}}$ follow a stochastic process, i.e. that they form a Markov chain $\Pr(X^{n+1} = x^{n+1} | X^n = x^n) = \Pr(X^{n+1} = x^{n+1} | X^n = x^n, \dots, X^0 = x^0)$ and obey the recurrence relation $\rho^t = G^t \rho$ for some stochastic map G over configurations.

Our problem is to determine what it means for a stochastic map over configurations G to be *causal*, meaning that arbitrarily remote regions I and J do not influence each other *by any means*. Then PCA will just be causal, shift-invariant stochastic maps.

3 Continuity

In the deterministic case CA were axiomatized by the celebrated Curtis-Lyndon-Hedlund Theorem[13], which we now recall. First the space of configurations is endowed with a metric:

Definition 4 (Metric) The function $d(\cdot, \cdot) : \mathcal{C} \times \mathcal{C} \rightarrow \mathbb{R}^+$ such that $d(c, c') = 0$ if $c = c'$ and $d(c, c') = 1/2^k$ with $k = \min\{i \in \mathbb{N} \mid c_{-i \dots i} \neq c'_{-i \dots i}\}$ is a metric. For $c, c' \in \mathcal{C}$ and $\varepsilon > 0$ we have (with $n = \lfloor \log_2(\varepsilon) \rfloor$): $d(c, c') < \varepsilon \Leftrightarrow c_{-n \dots n} = c'_{-n \dots n}$.

Definition 5 ((Uniform) continuity) A function $F : \mathcal{C} \rightarrow \mathcal{C}$ is continuous if and only if for all $c \in \mathcal{C}$ and $\varepsilon > 0$, there exists $\eta > 0$ such that for all $c' \in \mathcal{C}$: $d(c, c') < \eta \Rightarrow d(F(c), F(c')) < \varepsilon$. A function $F : \mathcal{C} \rightarrow \mathcal{C}$ is uniformly continuous if and only if for all $\varepsilon > 0$, there exists $\eta > 0$ such that for all $c, c' \in \mathcal{C}$: $d(c, c') < \eta \Rightarrow d(F(c), F(c')) < \varepsilon$.

In other words a function $F : \mathcal{C} \longrightarrow \mathcal{C}$ is uniformly continuous if and only if for all $n \in \mathbb{N}$, there exists $m \in \mathbb{N}$ such that for all $c, c' \in \mathcal{C}_\infty$, $c_{-m \dots m} = c'_{-m \dots m}$ implies $F(c)_{-n \dots n} = F(c')_{-n \dots n}$. Notice that rephrased in this manner, uniform continuity is a synonym for non-signalling, i.e. the fact that information does not propagate faster than a fixed speed bound. Continuity on the other hand expresses a somewhat strange form of relaxed non-signalling, where information does not propagate faster than a certain speed bound, but this speed bound depends upon the input. However, it so happens that the two notions coincide for compact spaces. Moreover, classical CA are easily defined upon infinite configurations $\mathcal{C}_\infty : \mathbb{Z} \rightarrow \Sigma$, for which the same $d(\cdot, \cdot)$ happens to be a compact metric. This yields:

Theorem 1 (Curtis, Lyndon, Hedlund)

A function $F : \mathcal{C}_\infty \longrightarrow \mathcal{C}_\infty$ is continuous and shift-invariant if and only if it is the global evolution of a cellular automaton over \mathcal{C}_∞ .

In other words this theorem just states that CA are exactly the non-signalling, shift-invariant functions. But instead of having to call them ‘non-signalling’ (a.k.a ‘uniformly continuous’), it only needs to call them ‘continuous’, due to the peculiarities of \mathcal{C}_∞ . However for the finite, yet unbounded configurations \mathcal{C} , the metric $d(\cdot, \cdot)$ is not compact. In this case we must assume the stronger, uniform continuity for the theorem to work. Generally speaking, it is rather difficult to find a relevant compact metric for probabilistic extensions of \mathcal{C}_∞ — and not worth the effort for the sole purpose of axiomatizing PCA. Indeed, let us directly assume the probabilistic counterpart of non-signalling (a.k.a uniform continuity for some extended metric):

Definition 6 (Non-signalling) *A stochastic map over configurations G is non-signalling if and only if for any ρ, ρ' two states over configurations, and for any cell i , we have:*

$$\rho_{i-1,i} = \rho'_{i-1,i} \quad \Rightarrow \quad (G\rho)_i = (G\rho')_i.$$

For example, *Parity* (see Figure 1) is non-signalling by construction. Is it reasonable to say, à la Hedlund, that PCA are the non-signalling, shift-invariant stochastic maps? Surprisingly, this is not the case. Imagine that Alice in Paris (cell 0) tosses a fair coin, whilst Bob in New York (cell $n+1$) tosses another. Imagine that the two coins are magically correlated, i.e. it so happens that they always yield the same result. Such a stochastic map is clearly not implementable by local mechanisms: we definitely need some amount of (prior) communication between Alice and Bob in order to make it happen. Yet it can be, as in *Magic-Coin*s (see Figure 2), perfectly non-signalling. While the setup cannot be used to communicate ‘information’ between distant places, it can be used to create spontaneous ‘correlations’ between them. We must forbid this from happening. In this respect, assuming only (non-uniform) continuity is the wrong direction to take.

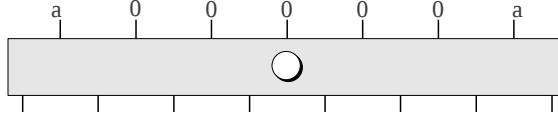


Fig. 2. *MagicCoins*. Inputs are ignored. Hence output random variables $(X_i)_{i=1\dots n}$ have the property that for any output x , the state (a.k.a the law of distribution) ρ_x is independent of the inputs. Yet output random variables X_1 and X_n are both uniformly distributed and maximally correlated. As such *MagicCoins* is a localized stochastic map, but can be modified into a translation-invariant stochastic map over entire configurations using the method described for *Parity*. *MagicCoins* is non-signalling, but it must not be considered a valid PCA: it is not non-correlating.

4 Avoiding spontaneous correlations

From the previous discussion we are obliged to conclude that the formalization of a robust notion of the causality of a dynamics is indeed a non-trivial matter in a probabilistic setting. From the *MagicCoins* example, we draw the conclusion that such a notion must forbid the creation of spontaneous correlations between distant places. The following definition clarifies what is meant by (non-)correlation between subsystems.

Definition 7 (Independence, tensor, trace) *Let I and J be disjoint. Stating that X_I and X_J are independent is equivalent to stating that for any uv in $(I \cup J \rightarrow \Sigma)$, $\rho_{I \cup J}(uv) = \rho_I(u)\rho_J(v)$, with u (resp. v) the restriction of uv to I (resp. J). In this case we write $\rho_{I \cup J} = \rho_I \otimes \rho_J$. This notation is justified because whenever I and J are finite, we have that the law of distribution ($\rho_{I \cup J}(uv)$) equals $(\rho_I(u)\rho_J(v))$ which is the definition of $(\rho_I(u)) \otimes (\rho_J(v))$, where \otimes is the Kronecker/tensor product. Whether or not X_I and X_J are independent, we can always recover ρ_I as the marginal of $\rho_{I \cup J}$ by averaging over every possible v , an operation which we denote Tr_J and call the trace-out/marginal-out operation. Namely we have that $\rho_I = Tr_J(\rho_{I \cup J})$, with $Tr_J(\rho_{I \cup J})(u) = \sum_v \rho_{I \cup J}(uv)$.*

A way to forbid spontaneous correlations is to require that, after one-time step of the global evolution G applied upon any initially fully determined configuration \tilde{c} , and for any two distant regions $I = -\infty \dots x-1$ and $J = x+1 \dots \infty$, the output $\rho = G\tilde{c}$ be such that $\rho_{I \cup J} = \rho_I \otimes \rho_J$. In other words remote regions remain independent. This formulation is somewhat cumbersome, because it seeks to capture in the vocabulary of ‘states’ a property which really belongs to their ‘dynamics’. The following definition clarifies what it means for a stochastic map to be localized upon a subsystem.

Definition 8 (Extension, localization, tensor) *Let I and J be disjoint. Consider a stochastic map S over I , i.e. from states $((I \rightarrow \Sigma) \rightarrow [0,1])$ to themselves. Then S can be trivially extended into a stochastic map $S \otimes Id$ over $I \cup J$, i.e. from states $((I \cup J \rightarrow \Sigma) \rightarrow [0,1])$ to themselves, and such that $(S \otimes Id)\rho_I \otimes \rho_J = (S\rho_I) \otimes \rho_J$. Whenever I and J are finite, this extends the*

stochastic matrix $(S\tilde{u}(u'))$ into $(S\tilde{u}(u')\delta_{vv'})$ of $(S \otimes Id)$. We say that some stochastic map over $I \cup J$ is localized upon I precisely if it arises in such away, i.e. as the trivial extension of some stochastic map S over I . Moreover if T is over J , then $(S \otimes T) = (Id \otimes T)(S \otimes Id) = (S \otimes Id)(Id \otimes T)$. This notation is justified because whenever I and J are finite, we have that the resulting stochastic matrix is $(S\tilde{u}(u')T\tilde{v}(v'))$ which equals $(S\tilde{u}(u')) \otimes (T\tilde{v}(v'))$, the Kronecker/tensor product of both stochastic matrices.

We can then forbid spontaneous correlations directly in terms of dynamics:

Definition 9 (non-correlation) A stochastic map over configurations G is non-correlating if and only if for any output cell i , there exist stochastic maps A, B acting over input cells $-\infty \dots x - 1$ and $x \dots +\infty$ respectively, such that: $Tr_x \circ G = A \otimes B$.

For example, *Parity* (see Figure 1) is non-correlating by construction. Now, is it reasonable to say that PCA are the non-correlating, shift-invariant stochastic maps? Amazingly, this is not the case. Indeed, consider a small variation of *Parity*, which we call *GenNLBox* and define in Figure 3. Such a stochastic map is clearly not implementable by local mechanisms: we definitely need some amount of communication between Alice and Bob in order to make it happen. It suffices to notice that *GenNLBox* is in fact a generalization of the ‘non-local box’, which we recover for $n = 2$, see Figure 3(a). But then, the non-local box owes its name precisely to the fact that it is not implementable by local mechanisms. Formal proofs of this assertion can be found in the literature [4, 21] and rely on the fact that the *NLBox* (maximally) violates the Bell inequalities [5, 30]. Yet *GenNLBox* (see Figure 3 (a)), was perfectly non-correlating. Hence, whilst the set-up cannot be used to communicate ‘information’ between distant places (it is non-signalling), nor to create spontaneous ‘correlations’ between distant places (it is non-correlating), we still must forbid it from happening!

5 Common cause principle

From the *GenNLBox* example of the previous Section, we are obliged to conclude that a robust notion of the causality of a dynamics in a probabilistic setting cannot be phrased just in terms of a non-signalling or a non-correlation property. Yet, this example has a virtue: it points towards similar questions raised in theoretical physics. Hence, this suggests looking at how those issues were addressed so far in theoretical physics.

Indeed, the *NLBox* is generally regarded as ‘unphysical’ because it does not comply with Bell’s [5] ‘local causality’, meaning that there is no prior shared resource that one may distribute between Alice and Bob (the outputs of the middle box of Figure 3 (b)), that will then allow them to perform the required task separately. Distributing quantum resources instead of classical resources (imagining that the outputs of the middle box can now be entangled quantum states) will not fix the problem: yes it does assist Alice and Bob in performing the required

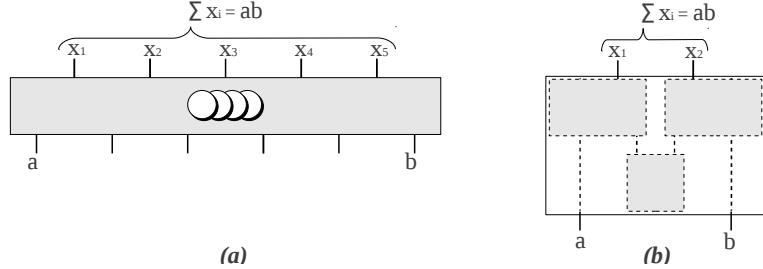


Fig. 3. (a) *GenNLBox*. Input cells $1 \dots n - 1$ are ignored. Output random variables $(X_i)_{i=1 \dots n}$ have the property that for strict subset I of $1 \dots n$, X_I is uniformly distributed, yet the global parity is always ab . As such *GenNLBox* is a localized stochastic map, but can be modified into a translation-invariant stochastic map over entire configurations using the method described for *Parity*. *GenNLBox* is both non-signalling and non-correlating, but it must not be considered a valid PCA: it is not V -causal. (b) *NLBox*. The output random variables X_1 and X_2 are uniformly distributed, but their parity is always equal ab . No circuit of the displayed dotted form can meet these specifications [5, 30].

task separately [2]; but only approximately so [28].

Bell's 'local causality' [5], 'Screening-off' [11], 'Stochastic Einstein locality' [14, 8], are all similar conditions, which stem from improvements upon Reichenbach's 'Principle of Common Cause' [16, 22], as was nicely reviewed in [15]. The common cause principle can be summarized as follows: "Two events can be correlated at a certain time if and only if, conditional upon the events of their mutual past, they become independent". In the context of this paper, this gives:

Definition 10 (Screening-off) *A stochastic map over configurations G obeys the screening-off condition if and only if for any input cell i with values in Σ , we have that there exists stochastic maps $(A_x, B_x)_{x \in \Sigma}$ acting over input cells $-\infty \dots i - 1$ and $i + 1 \dots +\infty$ respectively, such that for any ρ :*

$$\rho_i = \tilde{x} \quad \Rightarrow \quad G\rho = (A_x \otimes B_x)\rho.$$

Here input i is said to screen-off G .

This screening-off condition is physically motivated and does not suffer the problems that the non-signalling and non-correlation conditions had. Unfortunately however, it suffers most of the problems of the original, Standard-PCA definition: it is again incomplete and non-composable. For example, *Parity* does not obey the screening-off condition. Yet, *Parity* is a natural PCA, clearly implementable by local mechanisms as was shown in Figure 1. But the prior shared resource which is necessary in order to separate Alice from Bob is not present in the input cells, rather it is generated on-the-fly within one time-step of G , c.f. the circle-marked boxes. In other words, the reason why the screening-off condition rejects *Parity*, is because the condition is too stringent in demanding that

screening-off events be made explicit in the inputs x . A more relaxed condition would be to require that x may be completed so as to then screen-off G .

Definition 11 (Screening-off-completable) *A stochastic map over configurations G is screening-off-completable, or simply V -causal if and only if for any input cell i , we have that there exists stochastic maps (with input/output ranges marked as sub/superscript indices) $A_{-\infty \dots i-1}^{l'}, L_{l',l}^i, C_i^{l,r}, R_{r,r'}^{i+1 \dots \infty}$ such that:*

$$G = (L \otimes R)(A \otimes C \otimes B).$$

Here C is said to screen-off G at i . See Figure 4 (a).

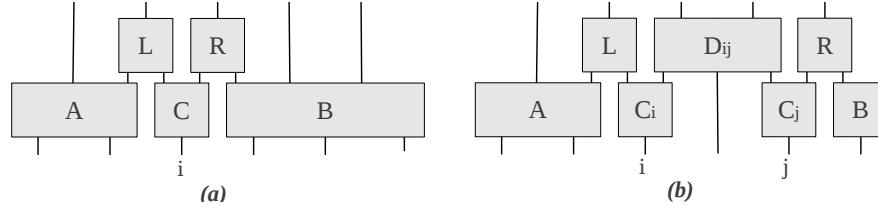


Fig. 4. (a) G is V -causal if it can be put in this form, for each i . (b) A strengthened condition: VV -causality.

Again, is it reasonable to say that PCA are the screening-off-completable, shift-invariant stochastic maps? Again, this is not the case. Indeed, consider a small variation of *Parity*, which we call *VBox* and define in Figure 5. Such a stochastic map is not implementable by local mechanisms: we need some amount of communication between Alice and Bob in order to make it happen. Yet *VBox* is perfectly V -causal, as shown in Figure 5. Hence, whilst the set-up is screening-off-completable, we still must forbid it, our condition is again too weak. A natural family of conditions to consider is VV -causality (as defined in Figure 4 (b)), V^3 -causality etc. This route seems a dead end; we believe that the *VBox* is the $k = 1$ instance of the following more general result:

Conjecture 1 *For all k , there exists a V^k -Box which is V^k -causal but not V^{k+1} -causal.*

6 Concluding definition

Our best definition. We have examined several, well-motivated causality principles (non-signalling, non-correlation, V -causality) and shown, through a series of surprising counter-examples, that stochastic maps with these properties are not necessarily implementable by local mechanisms. In the limit when k goes to infinity, V^k -causality turns into the following definition (assuming shift-invariance):

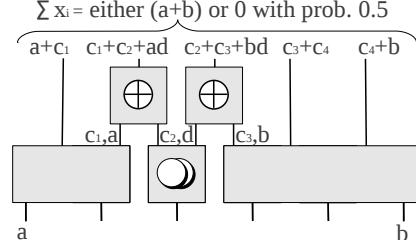


Fig. 5. *VBox*. Input cells $1 \dots n - 1$ are ignored. Output random variables $(X_i)_{i=1 \dots n}$ have the property that for strict subset I of $1 \dots n$, X_I is uniformly distributed, yet their global parity is, with equal probability, ab or 0. As such *VBox* is a localized stochastic map, but can be modified into a translation-invariant stochastic map over entire configurations using the method described for *Parity*. *VBox* is non-signalling, non-correlating and *V*-causal, but it must not be considered a valid PCA: it is not *VV*-causal. The *VBox* can be seen as a chain of two *GenNLBoxes*, chaining more of them yields a V^k *Box*.

Definition 12 (Probabilistic Cellular Automata) *A stochastic map over configurations G is a PCA if and only if*

$$G = (\bigotimes D)(\bigotimes C)$$

where the i^{th} stochastic matrix C has input i and outputs l_i, r_i , and the i^{th} stochastic matrix D has inputs r_{i-1}, l_i and output i , see Figure 6.

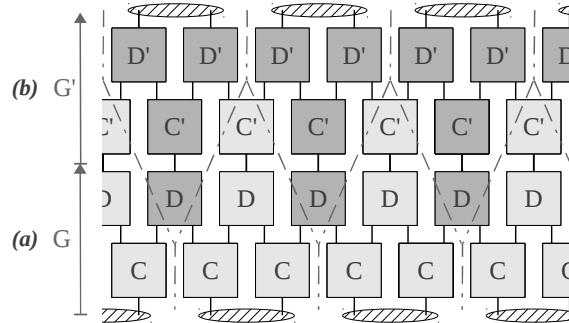


Fig. 6. (a) An operational definition of PCA. Time flows upwards. Wires stand for cells. The first two layers of boxes describe one application of a PCA $G = (\bigotimes D)(\bigotimes C)$. (b) Its compositionality. The second two layers describe one application of a PCA $G' = (\bigotimes D')(\bigotimes C')$. The resulting composition is a PCA $G'G = (\bigotimes(D' \otimes D')C'D)(\bigotimes C'D(C \otimes C))$ over the hatched supercells of size 2.

This definition has several advantages over Standard-PCA. First, it is complete, because it captures exactly what is meant (up to grouping) by a shift-invariant

stochastic map implementable by local mechanisms. Second, it is composable, as shown in Figure 6. Third, it is less ad-hoc. Indeed, we have mentioned in the introduction that much mathematics is done by having axiomatic definitions and operational description to coincide (usually by means of structure versus characterization theorems). It could be said that the same effect has been achieved in this paper, but in a different manner. Indeed, in this paper we have considered the natural candidate axiomatic definitions of PCA, whose limit is the natural operational description of PCA, and we have discarded each of them — by means of counter-examples. In this sense, we have pushed this family of candidate axiomatic definitions to its limit, i.e. as far as to coincide with the operational description.

Discussion. Still, we cannot really pretend to have reached an axiomatization, nor a characterization of PCA: Definition 12 can be presented as just the square of two Standard-PCA; or even more simply as a Standard-PCA with its two phases reversed: first a model of noise gets applied (a stochastic matrix is applied homogeneously on each individual cell), and second, a classical CA is applied. If anything, this paper has shown that it is rather unlikely that an axiomatization can be achieved. The authors are not in the habit of publishing negative results. However, the question of an characterization of PCA à la Hedlund has been a long-standing issue: we suspect that many researchers have attempted to obtain such a result in vain. At some point, it becomes just as important to discard the possibility of a Theorem as to establish a new one. Moreover, advances on the issue of causality principles [5, 8, 14, 16, 22], have mainly arisen from the discovery of counter-examples (See [15] and the more recent [10]), to which this paper adds a number. An impressive amount of literature [21, 3] focusses on the *NLBox* counter-example (as regards its comparison with quantum information processing, but also its own information processing power) and raise the question of n -party extensions: the *GenNLBox*, the *VBox*, and the *V^kBoxes* could prove useful in this respect.

Finally, let us point out that the observation that the operational description of PCA seems to admit no other axiomatization than itself is in sharp contrast with both the classical case (See Hedlund’s theorem Section 1), the reversible case (See [17]) and the quantum case (See [1]); for which the non-signalling condition alone suffices to entail localizability, i.e. implementation by local mechanisms. It is the moment when probabilities come into the picture (whether via stochastic maps or via quantum operations — this paper cancels out all efforts to generalize the result of [1] to an open systems setting) that the non-signalling condition becomes too weak. Moreover, we have seen that replacement principles based on a ‘common cause’ are hardly satisfactory. In philosophy of science it has been argued that the very notion of ‘physical law’ requires an underlying notion of causality, such as non-signalling. But probabilistic/stochastic theories seem to require a very explicit notion of causality, such as localizability.

Acknowledgements

The authors would like to acknowledge enlightening discussions with Joe Henson, Jean Mairesse, Jacques Mazoyer, Nicolas Schabanel, Rafaël Sorkin and Reinhard Werner.

References

1. P. Arrighi, V. Nesme, and R. Werner. Unitarity plus causality implies localizability. *QIP 2010 (long talk)*, ArXiv preprint: arXiv:0711.3975, 2010.
2. A. Aspect, J. Dalibard, and G. Roger. Experimental test of Bell’s inequalities using time-varying analyzers. *Physical Review Letters*, 49(25):1804–1807, 1982.
3. J. Barrett, L. Hardy, and K. Kent. No Signalling and Quantum Key Distribution. *Phys. Rev. Lett.*, 95(010503), 2005.
4. D. Beckman, D. Gottesman, M. A. Nielsen, and J. Preskill. Causal and localizable quantum operations. *Phys. Rev. A*, 64(052309), 2001.
5. J.S. Bell. On the Einstein Podolsky Rosen paradox. *Physics*, 1:195, 1964.
6. C. Bruun. *A model of consumption behaviour using cellular automata*. Aalborg University, 1996.
7. Ana Busic, Jean Mairesse, and Irene Marcovici. Probabilistic cellular automata, invariant measures, and perfect sampling. Pre-print arXiv:1010.3133, 2010.
8. J. Butterfield. Stochastic Einstein locality revisited. *The British Journal for the Philosophy of Science*, 58(4):805, 2007.
9. B. Chopard and M. Droz. *Cellular automata modeling of physical systems*. Cambridge University Press New York, 1998.
10. B. Coecke and R. Lal. Time-asymmetry and causal structure. Arxiv preprint arXiv:1010.4572, 2010.
11. D. Craig, F. Dowker, J. Henson, S. Major, D. Rideout, and R.D. Sorkin. A Bell inequality analog in quantum measure theory. *Journal of Physics A: Mathematical and Theoretical*, 40:501, 2007.
12. N. Fates. Stochastic Cellular Automata Solve the Density Classification Problem with an Arbitrary Precision. In *Proceedings of STACS*, to appear in *Lecture Notes in Computer Science*. Springer, 2011.
13. G. A. Hedlund. Endomorphisms and automorphisms of the shift dynamical system. *Math. Systems Theory*, 3:320–375, 1969.
14. G. Hellman. Stochastic Einstein-locality and the Bell theorems. *Synthese*, 53(3):461–503, 1982.
15. J. Henson. Comparing causality principles. *Studies In History and Philosophy of Science Part B: Studies In History and Philosophy of Modern Physics*, 36(3):519–543, 2005.
16. G. Hofer-Szabo, M. Redei, and L.E. Szabo. On Reichenbach’s common cause principle and Reichenbach’s notion of common cause. *The British Journal for the Philosophy of Science*, 50(3):377, 1999.
17. J. Kari. On the circuit depth of structurally reversible cellular automata. *Fundamenta Informaticae*, 38(1-2):93–107, 1999.
18. N. Margolus. Physics-like models of computation. *Physica D: Nonlinear Phenomena*, 10(1-2), 1984.
19. E. F. Moore. The firing squad synchronization problem. In *Sequential Machines, Selected Papers*, pages 213–214. Addison-Wesley, 1964.

20. K. Nagel and M. Schreckenberg. A cellular automaton model for freeway traffic. *J. Phys. I France*, 2:2221–2229, 1992.
21. S. Popescu and D. Rohrlich. Quantum nonlocality as an axiom. *Foundations of Physics*, 24(3):379–385, 1994.
22. M. Rédei and S.J. Summers. Local primitive causality and the common cause principle in quantum field theory. *Foundations of Physics*, 32(3):335–355, 2002.
23. D. Regnault, N. Schabanel, and É. Thierry. Progresses in the analysis of stochastic 2D cellular automata: a study of asynchronous 2D minority. *Theoretical Computer Science*, 410(47-49):4844–4855, 2009.
24. D.H. Rothman and S. Zaleski. *Lattice-gas cellular automata: simple models of complex hydrodynamics*. Cambridge University Press, 1997.
25. T. Toffoli and N. Margolus. *Cellular Automata Machine — A new Environment for Modelling*. MIT Press, Cambridge MA, 1987.
26. A. Toom. Cellular automata with errors: Problems for students of probability. *Topics in Contemporary Probability and Its Applications*, pages 117–157, 1995.
27. A. Toom, N. Vasilyev, O. Stavskaya, L. Mityushin, G. Kurdyumov, and S. Pirogov. Discrete Local Markov Systems. Stochastic Cellular Systems: ergodicity, memory, morphogenesis. Ed. by R. Dobrushin, V. Kryukov and A. Toom. Nonlinear Science: theory and applications, 1990.
28. B.S. Tsirelson. *Lett. Math. Phys.*, 4(93), 1980.
29. J. von Neumann. *Theory of Self-Reproducing Automata*. University of Illinois Press, Champaign, IL, USA, 1966.
30. R.F. Werner and M.M. Wolf. All-multipartite Bell-correlation inequalities for two dichotomic observables per site. *Physical Review A*, 64(3):32112, 2001.
31. R. White and G. Engelen. Cellular automata as the basis of integrated dynamic regional modelling. *Environment and Planning B*, 24:235–246, 1997.
32. D. A. Wolf-Gladrow. Lattice-Gas Cellular Automata and Lattice Boltzmann Models. In *Lecture Notes in Mathematics*. Springer, 2000.
33. K. Zuse. Rechnender Raum. Elektronische Datenverarbeitung. *English Translation: Calculating Space, MIT Tech. Translation*, 8:336–344, 1967.